

TECHNIQUE FOR ADMISSION CONTROL OF PACKET FLOWS

CROSS-REFERENCE TO RELATED APPLICATIONS

5 This patent application claims priority to U.S. Provisional
Patent Application No. 60/497,932 (Attorney Docket No.
57983.000156; Client Reference No. 16347ROUS01P), filed August
27, 2003, which is hereby incorporated by reference herein in
its entirety.

FIELD OF THE INVENTION

10 The present invention relates generally to computer and
communications network and, more particularly, to a technique
for admission control of packet flows.

BACKGROUND OF THE INVENTION

15 In recent years, voice over IP (VoIP) has become a new
standard for voice communications due to the ubiquity of
Internet Protocol (IP) networks. Compared with traditional
20 telephonic technology, VoIP offers the benefits of cost savings,
open standards, multi-vendor interoperability, and integrated
voice and data networks. Other real-time applications include,
for example, video-conferencing, multimedia streaming and other
multimedia services that require guaranteed quality of service
25 (QoS), all of which may benefit from the technique for admission

control of packet flows in accordance with the present invention.

Compared with traditional telephonic technology, VoIP also has its own difficulties. Traditionally voice communications have been carried over the public switched telephone network (PSTN), a circuit-switched network that can ensure the quality of each call through dedicated bandwidth allocations. This consistently high level of voice quality is called "toll quality." Unlike PSTN, IP networks are based on packet switching rather than circuit switching. Packet switching works well for data communications because it maximizes bandwidth utilization by allowing all users to dynamically share network bandwidth. However, traditional packet switching techniques do not adequately recognize the mixing of many different applications that have different traffic flow characteristics and performance requirements in one network or link. For example, real-time applications typically require minimum delay, low packet loss (or guaranteed bandwidth), whereas non-real-time application may be able to tolerate longer delays and recover from packet loss. Though telephony service, video-conferencing and one-way multimedia streaming are real-time applications, the last two are more bandwidth-demanding.

Some providers employ bandwidth management techniques such as prioritization to ensure that VoIP and other real-time multimedia applications get what they need. However, bandwidth management alone simply allocates bandwidth to critical

5 applications at the expense of other applications and does not limit or prevent real-time applications from congesting a link to a point where the performance or quality of service offered for all real-time flows become unacceptable. Some providers or enterprises over-provision their network capacity so that
10 application demands would never hit the bandwidth limit. This solution is not economical since it may require expensive upgrades to the carrier networks to meet the increasing peak demands for bandwidth. None of the existing solutions can prevent network congestion in an efficient and cost-effective
15 manner and still provide high quality of service for real-time flows.

In view of the foregoing, it would be desirable to provide a solution which overcomes the above-described inadequacies and shortcomings.

20

SUMMARY OF THE INVENTION

According to the present invention, a technique for admission control of packet flows is provided. In one

particular exemplary embodiment, the technique may be realized as a method for admission control of packet flows in a network.

The method may comprise determining at least one flow rate associated with a plurality of packets. The method may also

5 comprise marking at least one predetermined bit in at least one of the plurality of packets if the at least one flow rate is greater than a predetermined rate. The method may further comprise controlling an admission of additional packets into the network based at least in part on the marking of the at least
10 one predetermined bit in the at least one of the plurality of packets.

In accordance with other aspects of this particular exemplary embodiment of the present invention, the network may comprise a plurality of network elements, and the at least one
15 flow rate may be determined at a first network element, where the first network element is part of an access link of the network.

In accordance with further aspects of this particular exemplary embodiment of the present invention, the at least one
20 of the plurality of packets may comprise at least one signaling packet. The at least one signaling packet may originate from an end terminal outside the network. Information associated with the at least one predetermined bit in the at least one signaling

packet may be communicated to the end terminal. And the end terminal may echo information associated with the at least one predetermined bit in the at least one signaling packet in a transmission to the network.

5 In accordance with still further aspects of this particular exemplary embodiment of the present invention, the method may further comprise denying the admission of the additional packets into the network if the at least one predetermined bit in the at least one of the plurality of packets is marked.

10 In accordance with additional aspects of this particular exemplary embodiment of the present invention, the admission of the additional packets into the network may be controlled by an entity that controls the network. The admission of the additional packets may be based at least in part on priorities
15 or importance of the plurality of packets and the additional packets.

 In accordance with another aspect of this particular exemplary embodiment of the present invention, the plurality of packets may comprise real-time packets. The plurality of
20 packets may comprise Internet Protocol (IP) packets, including voice over IP (VoIP) packets. And the at least one predetermined bit may be part of a Differentiated Services field in an IP header of the at least one of the plurality of packets.

In accordance with yet another aspect of this particular exemplary embodiment of the present invention, the predetermined rate may be based on a network bandwidth allocated for the plurality of packets. And the predetermined rate may be raised
5 to a value above the allocated network bandwidth for a predetermined period of time.

In another particular exemplary embodiment, the technique may be realized by at least one signal embodied in at least one carrier wave for transmitting a computer program of instructions
10 configured to be readable by at least one processor for instructing the at least one processor to execute a computer process for performing the method as recited above.

In yet another particular exemplary embodiment, the technique may be realized by at least one processor readable
15 carrier for storing a computer program of instructions configured to be readable by at least one processor for instructing the at least one processor to execute a computer process for performing the method as recited above.

In still another particular exemplary embodiment, the
20 technique may be realized by a system for admission control of packet flows. The system may comprise duplicating at least one network element that determines at least one flow rate associated with a plurality of packets, and marks at least one

predetermined bit in at least one of the plurality of packets if the at least one flow rate is greater than a predetermined rate. The system may further comprise an admission control module that controls an admission of additional packets into the network

5 based at least in part on the marking of the at least one predetermined bit in the at least one of the plurality of packets.

In a further particular exemplary embodiment, the technique may be realized by a system for admission control of packet
10 flows. The system may comprise means for determining at least one flow rate associated with a plurality of packets. The system may also comprise means for marking at least one predetermined bit in at least one of the plurality of packets if the at least one flow rate is greater than a predetermined rate.
15 The system may further comprise means for controlling an admission of additional packets into the network based at least in part on the marking of the at least one predetermined bit in the at least one of the plurality of packets.

The present invention will now be described in more detail
20 with reference to exemplary embodiments thereof as shown in the accompanying drawings. While the present invention is described below with reference to exemplary embodiments, it should be understood that the present invention is not limited thereto.

Those of ordinary skill in the art having access to the teachings herein will recognize additional implementations, modifications, and embodiments, as well as other fields of use, which are within the scope of the present invention as disclosed
5 and claimed herein, and with respect to which the present invention could be of significant utility.

BRIEF DESCRIPTION OF THE DRAWINGS

In order to facilitate a fuller understanding of the
10 present invention, reference is now made to the accompanying drawings, in which like elements are referenced with like numerals. These drawings should not be construed as limiting the present invention, but are intended to be exemplary only.

Figure 1 is a flow chart illustrating an exemplary method
15 for admission control of packet flows in accordance with an embodiment of the present invention.

Figure 2 is a diagram illustrating the DiffServ Field in an IPv4 packet header in accordance with an embodiment of the present invention.

20 Figure 3 is a block diagram illustrating an exemplary system for admission control of packet flows in accordance with an embodiment of the present invention.

Figure 4 is an exemplary bandwidth diagram in accordance

with an embodiment of the present invention.

Figure 5 is a schematic illustration of an admission control scenario in accordance with an embodiment of the present invention.

5 Figure 6 is a schematic illustration of another admission control scenario in accordance with an embodiment of the present invention.

Figure 7 is a schematic illustration of yet another admission control scenario in accordance with an embodiment of
10 the present invention.

Figure 8 is a schematic illustration of still another admission control scenario in accordance with an embodiment of the present invention.

15 **DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENT(S)**

For illustration purposes, the technique for admission control of packet flows in accordance with the present invention will be described below with specific reference to voice over IP (VoIP) packet flows under current version of Internet Protocol
20 (IPv4). However, it should be appreciated that the technique is applicable to any real-time packet flows such as voice, video, data and combinations thereof. The technique is also applicable to future versions of Internet Protocol including IPv6.

Referring to Figure 1, there is shown a flow chart illustrating an exemplary method for admission control of packet flows in accordance with an embodiment of the present invention.

In step 100, a first end terminal (VoIP terminal at call
5 originating end) may initiate a call to a second end terminal
(VoIP terminal at call terminating end). The call setup may be
started by using a number of standardized and/or customized
signaling methods, such as Nortel Networks UniStim Interface
Protocol (UNISTIM), Media Gateway Control Protocol (MGCP),
10 Session Initiation Protocol (SIP), Automated System for
Performance Evaluation of Networks (ASPEN), and the
International Telecommunication Union (ITU) defined H.248 or
H.323 standards can be used.

In step 102, a network edge element may meter a flow of
15 VoIP packets entering the network (i.e., ingress flow). A
network edge element may be any network element, such as a
router, switch or middle-box, located on or near the edge of a
computer and/or communications network. The edge element may be
part of an access port or access link through which VoIP packets
20 flow into and/or out of the network.

VoIP packet flows typically include two types of packets:
signaling packets and payload packets. The signaling packets
carry call-control messages that are used to set up and manage

connections between VoIP terminals. The payload packets carry the digitally encoded voice conversations. According to one embodiment of the invention, the signaling packets and payload packets may be transmitted using different Differentiated

5 Services Code Points (DSCP). In IPv4, the Differentiated Services (DiffServ or DS) model defines different service class levels for different real-time packets. Figure 2 is a diagram illustrating the DiffServ Field in an IPv4 packet in accordance with an embodiment of the present invention. The six most
10 significant bits (MSBs) of the DiffServ Field, also known as Differentiated Services Code Point (DSCP), serve as an indication of the packet's required service class level. The two least significant bits (LSBs) are reserved as Explicit Congestion Notification (ECN) bits. According to an embodiment
15 of the invention, a specific service class may be used for VoIP traffic. For example, Expedited Forwarding (EF) DSCP may be used to mark VoIP payload and CS5 DSCP may be used for signaling. Packets marked with both DSCP values may be forwarded using the same service class.

20 A DiffServ style traffic conditioner, which is capable of metering VoIP flows and marking the ECN bit(s) of DS Field in IP headers, may be implemented on the network edge element. For example, a sum of EF and CS5 marked packets may be metered for a

predetermined period of time. As a result, a flow rate may be determined. The flow rate may be measured for ingress flows and/or egress flows. An ingress flow rate is an indication of the uplink bandwidth usage, and an egress flow rate is an indication of the downlink bandwidth usage.

In step 104, it may be determined whether the ingress flow rate is greater than a predetermined maximum rate "A." The maximum rate "A" may be a predetermined portion of a total bandwidth allocated for VoIP applications. For example, if an access link has 10 Mbps bandwidth, 1.5 Mbps of which is allocated for VoIP service, the maximum rate "A" may be 1.5 Mbps, with the remaining 8.5 Mbps allocated for other traffic. Emergency "911" calls may be allowed to borrow for short period of time bandwidth that was allocated for other traffic, therefore total of normal and emergency VoIP flows may exceed the 1.5 Mbps rate.

If the flow rate exceeds "A," the uplink is congested (uplink has reached the bandwidth limit for VoIP flows). Then the network edge element may, in step 106, mark the ECN bit(s) in the IP headers of the signaling packets flowing into the network. The signaling packets may be those sent by the first end terminal to set up new Internet calls. The signaling packets may be part of an ingress flow coming into the network.

According to embodiments of the invention, the ECN bit(s) may be encoded by the network edge elements to convey a status of bandwidth usage in the access link. For example, Bit 7 of the DiffServ Field may be set to indicate that the access link is congested with VoIP traffic. Or both ECN bits, i.e., Bit 6 and Bit 7, may be encoded to indicate levels of congestion in the access link.

According to embodiments of the invention, the method of marking or encoding ECN bit(s) may be extended to data fields other than the ECN bits or the DS Field. Depending on the desired application and the specific signaling method, other predetermined bit(s) in the signaling packets may be marked or encoded to convey the bandwidth usage status. Nor is this method limited to signaling packets. Other types of packets, such as control packets and payload packets may also serve the same or similar purpose. For example, the predetermined bits may be in the IP header, UDP header, RTP header, or payload header of the packet that is used to convey the bandwidth usage status.

In step 108, an admission control decision point may examine the ECN bit(s) in the signaling packets it receives. An admission control decision point may be a communication server or some other server in the network, for example. For each

signaling packet it receives, the admission control decision point in a communication server may determine, in step 110, whether the ECN bits have been marked. If so, the decision point may, in step 112, stop the call setup process and may also
5 send cause of failure to the VoIP terminal at call originating end.

If, in step 110, it is determined that the ECN bits of the signaling packets are not marked, call setup may continue in step 114. The communication server may next send message(s) to
10 the call originating terminal.

In step 116, the network edge element may meter a flow of VoIP packets exiting the network (i.e., egress flow). In step 118, it is determined whether the egress flow rate is greater than a predetermined maximum rate "B." If so, the downlink is
15 congested. Then, in step 120, the network edge element may mark the ECN bit(s) of DS Fields in the signaling message(s) which was sent from the communication server. If the egress flow rate is no more than the predetermined maximum, the ECN bit(s) in the signaling message(s) will not be marked.

20 In step 122, as the message(s) reach the call originating terminal, the ECN bit settings in the message(s) are echoed in the subsequent messages that are sent to the communication server. These messages will again pass through the network edge

element on their way to the communication server. As depicted in steps 102 through 106, the network edge element will mark the appropriate ECN bit(s) if it meters a VoIP flow rate greater than the predetermined maximum. And these ECN bits will be examined by the admission control decision point in the communication server to determine the bandwidth status on the access link. Whether the ECN bits have been marked in step 120 or in step 106, the admission control decision point will be notified and will be able to take appropriate measures to continue or to end the call setup. If the ECN bits are not marked, the communication server may continue the call setup process.

According to embodiments of the invention, new emergency calls may still be allowed even when normal calls are denied access to the network. In fact, a network implementing the method for admission control of packet flows may over-subscribe VoIP calls for a certain period of time without negative impact to an overall performance of the network. An exemplary bandwidth diagram illustrating temporary over-subscription is shown in Figure 4. In this exemplary diagram, a total bandwidth of the network is divided into two parts. One part is allocated for real-time traffic such as VoIP and video-telephony and is forwarded by a separate service or forwarding class. The rest

of the bandwidth is allocated for traffic in other service or forwarding classes. In practice, a combination of normal calls and emergency calls may be allowed to use more bandwidth than what is allocated for voice traffic, up to the policing

5 bandwidth. For example, up till time T_1 the network may have been servicing normal calls using all the bandwidth allocated for voice traffic. Then, a number of emergency calls may request setup. Instead of dropping ongoing normal calls to make bandwidth available for emergency calls, the network may stop
10 admitting additional normal calls and use a portion of the bandwidth from other service classes to service the incoming emergency calls. Within approximately three minutes, which is the average length of normal calls, a number of voice calls will terminate normally, freeing up significant portion of bandwidth
15 in the voice service class. After then, emergency calls may be allowed to take up most or all of the bandwidth allocated for voice service class.

Figure 3 is a block diagram illustrating an exemplary system for admission control of packet flows in accordance with
20 an embodiment of the present invention. The system may comprise a plurality of network edge elements (e.g., Router 304, Router 306, Router 310 and Router 312) located on or near the edge of an IP Network 300. The system may further comprise a

Communication Server 302 serving as an admission control
decision point in this particular embodiment. Routers 304, 306,
310 and 312 serve as part of a bearer (payload) connection path
328 between end terminals 308 and 314. There are also other end
5 terminals (316, 318, 320 and 322) that are served by these
routers. IP Network 300 is also connected to a public switched
telephone network (PSTN) 324 and a Time-Division Multiplexing
(TDM) network 326 through their respective media gateways. A
connection 330 between Routers 304 and 306 serves as an access
10 link for end terminals 308, 320 and 322. A connection 332
between Routers 310 and 312 serves as an access link for end
terminals 314 and 316.

In order to initiate a VoIP call to end terminal 308, end
terminal 314 may send signaling packets through access link 332
15 to Communication Server 302. Either Router 312 or Router 310 or
both may be monitoring packet flows into and out of IP Network
300. For example, if Router 310 meters a flow rate that is
greater than a predetermined maximum rate for regular calls, it
may mark the ECN bit(s) in the signaling packets sent by end
20 terminal 314. As Communication Server 302 receives the
signaling packets, it may examine their ECN bit(s). If it
identifies marked signaling packets from Router 310,
Communication Server 302 may stop admitting new calls, with the

exception of "911" and other emergency calls. Communication
Server 302 may also send call-failure notifications to the end
terminal 314. If sufficient bandwidth has been verified on the
access link 332, a similar verification process may take place
5 on access link 330 which is on the receiving side of the IP
Network 300.

As mentioned above, the method for admission control of
packet flows in accordance with the present invention may be
implemented to verify bandwidth availability for either an
10 uplink or a downlink on both the call originating end and the
call terminating end. For a better understanding of the
different scenarios, a few examples are described below in
connection with Figures 5 through 8.

The examples illustrated in Figures 5 through 8 are based
15 on one possible implementation of the Session Initiation
Protocol (SIP).

In each of Figures 5 through 8, there is schematically
shown an End Terminal 1 at the call originating end, an Edge
Router I, a Communication Server 500, an Edge Router E and an
20 End Terminal 2 at the call terminating end.

Figure 5 illustrates a scenario where the uplink is
congested on the call originating end. To initiate a call to
End Terminal 2, End Terminal 1 may transmit INVITE or call-setup

request message to Communication Server 500. The signaling packets may be marked with CS5 DSCP value. Edge Router I may be monitoring network ingress flows on its interface to the subscriber. If the network ingress (uplink) flow rate exceeds a predetermined maximum rate, Edge Router I may mark the ECN bit(s) in the signaling packets sent by End Terminal 1. When Communication Server 500 receives these signaling packets, it may examine the ECN bits. If the ECN bits are marked, Communication Server 500 is notified that there is insufficient bandwidth available on the uplink of the call originating end to admit normal calls. As a result, Communication Server 500 may stop the call setup process for normal (non emergency "911") calls. The call failure cause may be sent to End Terminal 1.

Figure 6 illustrates a scenario where the downlink is congested on the call originating end. Edge Router I may monitor not only ingress flows but also egress flows on its interface to the subscriber. If the ingress flow rate is below the predetermined maximum rate, the signaling packets may flow through without being marked. When these signaling packets have been processed by Communication Server 500, Server 500 may send 407 Proxy Authorization Request messages through Edge Router I back to End Terminal 1. If Edge Router A measures an egress flow rate that exceeds the predetermined maximum rate, it may

mark the ECN bits in 407 message. End Terminal 1 may echo the ECN bit settings in the 407 message it receives, so that the ACK message it subsequently sends will bear the same ECN bit settings. Upon examining these ACK message with marked ECN bits, Communication Server 500 is notified of the insufficient bandwidth on the downlink. Action may then be taken to stop call setup and the call failure cause may be sent to End Terminal 1.

Figure 7 illustrates a scenario where the uplink is congested on the call terminating end (End Terminal 2). The signaling packets sent by Voice Terminal 1 may travel to and from the Communication Server 500 without being marked by Edge Router I, which indicates that neither uplink nor downlink on the call originating end is congested. As call setup continues, the INVITE (Offer) message sent by End Terminal 1 may be forwarded to the call terminating end. Since the downlink is not congested, ECN bits are not set. End Terminal 2 may echo the ECN settings in the INVITE (Offer) messages to the NOTIFY messages it subsequently sends to Server 500. When the NOTIFY message passes through Edge Router E, their ECN bits may be marked if the ingress link of Edge Router E is congested. Upon examining the NOTIFY message, Communication Server 500 may become notified of the congestion on the call terminating end.

The call setup process may be discontinued and the failure cause may be sent to End Terminal 1.

Figure 8 illustrates a scenario where the downlink is congested on the call terminating end. The signaling packets sent by Voice Terminal 1 may travel to and from the Communication Server 500 without being marked by Edge Router I, which indicates that neither uplink nor downlink on the call originating end is congested. As call setup continues, the INVITE (Offer) message sent by End Terminal 1 may be forwarded to the call terminating end. Similar to Edge Router I, Edge Router E may monitor not only ingress flows but also egress flows on its interface to the subscriber. If the egress flows exceed a predetermined maximum rate, Edge Router B may mark the INVITE (Offer) message that are being forwarded towards End Terminal 2. End Terminal 2 may echo the ECN bit settings in a subsequent transmission of NOTIFY messages. When Communication Server 500 examines these subsequently transmitted NOTIFY messages, it may become notified of the congestion on the downlink. Then the call setup may be discontinued and End Terminal 1 may be notified of the failure cause.

At this point it should be noted that the technique for admission control of packet flows in accordance with the present invention as described above typically involves the processing

of input data and the generation of output data to some extent.

This input data processing and output data generation may be implemented in hardware or software. For example, specific electronic components may be employed in a computer and/or

5 communications network or similar or related circuitry for implementing the functions associated with admission control of packet flows in accordance with the present invention as described above. Alternatively, one or more processors operating in accordance with stored instructions may implement
10 the functions associated with admission control of packet flows in accordance with the present invention as described above. If such is the case, it is within the scope of the present invention that such instructions may be stored on one or more processor readable carriers (e.g., a magnetic disk), or
15 transmitted to one or more processors via one or more signals.

The present invention is not to be limited in scope by the specific embodiments described herein. Indeed, other various embodiments of and modifications to the present invention, in addition to those described herein, will be apparent to those of
20 ordinary skill in the art from the foregoing description and accompanying drawings. Thus, such other embodiments and modifications are intended to fall within the scope of the following appended claims. Further, although the present

invention has been described herein in the context of a particular implementation in a particular environment for a particular purpose, those of ordinary skill in the art will recognize that its usefulness is not limited thereto and that

5 the present invention can be beneficially implemented in any number of environments for any number of purposes. Accordingly, the claims set forth below should be construed in view of the full breath and spirit of the present invention as disclosed herein.